

6

Neuroethology of Decision Making

Geoffrey K. Adams, Karli K. Watson,
John Pearson, and Michael L. Platt

Abstract

A neuroethological approach to decision making posits that neural circuits mediating choice evolved through natural selection to link sensory systems flexibly to motor output in a way that enhances the fit between behavior and the local environment. This chapter discusses basic prerequisites for a variety of decision systems from this viewpoint, with a focus on two of the best studied and most widely represented decision problems. The first is patch leaving, a prototype of environmentally based switching between action patterns. The second is social information seeking, a behavior that, while functionally distinct from foraging, can be addressed in a similar framework. It is argued that while the specific neural solutions to these problems sometimes differ across species, both the problems themselves and the algorithms instantiated by biological hardware are repeated widely throughout nature. The behavioral and mathematical study of ubiquitous decision processes like patch leaving and information seeking thus provides a powerful new approach to uncovering the fundamental design structure of nervous systems.

Introduction

In the wetlands of North America, male red-winged blackbirds (*Agelaius phoeniceus*) compete for territorial control over patchy resources, with some males managing to dominate rich territories while others are stuck with low-quality leavings. When the mating season begins, female red-winged blackbirds select a primary male with which to mate and subsequently nest in his territory. Because territory quality is so variable, some highly successful males may attract a harem of up to about fifteen females, while many unsuccessful males will fail to attract a mate at all. Although each female may produce multiple broods during the breeding season, there is a significant chance that she will fail to bring any offspring to fledging. With the stakes so high, her selection of a mate is a critical decision that will have direct consequences for her fitness.

Many factors can impact her success and may, in principle, play a role in her decision: these include male health and parasite load, food density, the presence of acceptable nesting sites, the number of females already present in the harem, the likelihood of receiving paternal care for her brood, and the presence of healthy male neighbors who may provide opportunities for extra-pair copulations. Despite decades of excellent studies of the red-winged blackbird, the extent to which females assess these factors in choosing a mate remains an active area of research (Beletsky 1996).

Mate choice in female red-winged blackbirds is an illustrative example of a complex decision-making problem posed by the particular details of an animal's habitat and social structure. Similar examples are common across a broad diversity of animal taxa: primates select long-term social partners with consequences for their health and reproductive success (Silk et al. 2003; Schulke et al. 2010; Silk et al. 2010); cleaner fish in coral reefs decide between providing good service or "cheating" their clients, and adjust their level of service for each individual client (Bshary and Noe 2003); African buffalo decide the best direction for the herd to move by implementing a vector average of the orientation favored by each cow weighted by her social status (Prins 1996); tens of thousands of honeybees in a swarm select a new nest site by comparing their own assessment of sites based on location, temperature, and humidity with the evaluations made and communicated by other scouts (Vonfrisch and Lindauer 1956; Seeley and Visscher 2004). In each case, the animal's immediate context may reasonably permit multiple possible behaviors, but only one or a few will optimize fitness.

Evolutionary theory and behavioral ecology identify the decision-making problems that animals face in their natural environments, as well as the costs, benefits, and constraints associated with pursuing specific behavioral strategies. In the case of mate choice in female red-winged blackbirds, one of the more successful proposals is the *polygyny threshold hypothesis*, which suggests that territorial resources are the primary factor influencing female choice (Verner and Willson 1966; Orians 1969). According to this model, females choose a mate based on the amount of resources his territory will offer her, accounting for the resources that will be consumed by the male's existing harem. Although other factors may also influence female choice, behavioral studies of red-winged blackbirds in their natural habitat suggest that the polygyny threshold hypothesis provides a good approximation of the true mate choice strategy that females pursue (Beletsky 1996).

Decision neuroethology is concerned with understanding the physiological mechanisms that evolution has produced for solving the complex decision-making problems posed by animals' environments. Although the study of animal behavior has revealed a remarkable diversity of such problems and their solutions, mathematical analysis has demonstrated that dissimilar-seeming problems may be solved with similar strategies; for example, patch and prey foraging (Stephens and Krebs 1986:32). When the fitness impact of a decision

is large, there are strong selective pressures on the corresponding decision-making mechanism to behave according to the mathematically optimal strategy. Such a constraint means that the neurobiological mechanisms mediating decisions may tend to be highly conserved or convergent across taxa for a general class of decision-making problems. To understand how decision-making behaviors evolve, the concepts of conservation, convergence, and disparity of mechanisms must be considered.

To illustrate the basic approach of decision neuroethology, we will describe a computational framework for understanding the selective pressures on decision-making mechanisms, and how to approach comparison between species. We will then discuss in greater detail two examples of decision problems in behavioral ecology and their neural implementations. First, we examine patch leaving, a prototype of environmentally based switching between action patterns. Second, we examine social information-seeking behaviors in nonhuman primates. We conclude that this work portends a more general understanding of complex decision problems and, ultimately, endorses the unification of theoretical and experimental work in behavioral ecology and neuroscience.

A Computational, Comparative Approach

In his seminal work on computational vision, Marr (1982) introduced a three-level framework for considering vision as a computational problem and argued that visual neuroscience must consider all three levels. Such a framework is also valuable in considering decision-making problems that must be solved by the nervous system (Figure 6.1). Furthermore, as we will see, this framework permits both a comparative analysis across species and across decision types within a single species.

According to Marr, the *computational level* of analysis is the level at which we consider the actual *problem* a biological system is attempting to solve. From an evolutionary standpoint, this is the level at which selective pressures exist to shape a decision process. Put another way, this is the ecological problem posed by an animal's niche. Of course, in an ultimate sense, this generally reduces to "maximize inclusive fitness." However, given that inclusive fitness may be practically impossible to measure in most contexts, we will often limit analysis to more specific goals such as "maximize resources for offspring" or "maximize long-term caloric intake." In the case of mate choice in the female red-winged blackbird, this level of analysis corresponds to the identification of factors associated with selecting a male that will influence seasonal reproductive success. In some well-known cases, problems posed at the computational level possess optimal solutions, such as the ones obtained by *marginal value theorem* (Charnov 1976; see also patch-leaving decisions discussed below). In this sense, an optimal solution is a relationship between environmental and behavioral parameters guaranteed to produce the best outcome under the posed

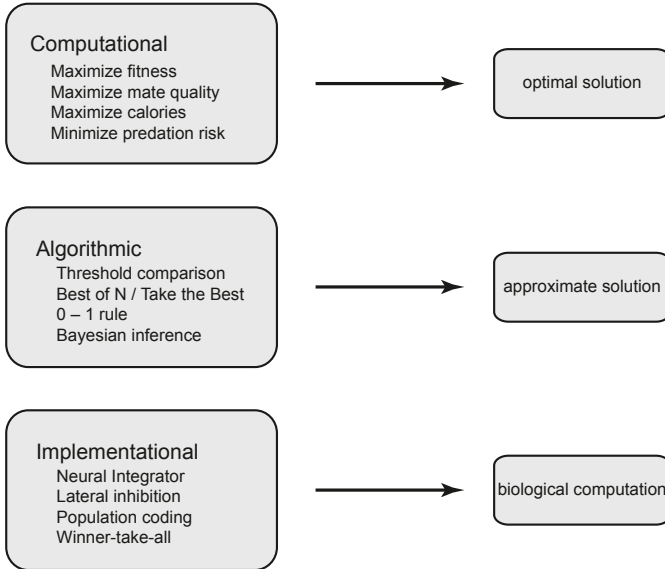


Figure 6.1 Decision-making problems can be analyzed at three levels. The computational level identifies the decision-making problem to be solved; the algorithmic level identifies strategies by which the problem may be solved; and the implementational level identifies the exact biological mechanism by which an organism solves the problem.

problem. There are many examples of organisms producing optimal or near-optimal behavior for solving an identified problem.

At the *algorithmic* and *representational level* of analysis, we are concerned with identifying *representations* (i.e., formal systems for describing features of the environment relevant to the problem) and *algorithms* (i.e., the series of computational steps whereby animals solve the problems posed). A popular choice of representation in neuroeconomic models of decision making is utility, but other representations are possible. The polygyny threshold hypothesis suggests a possible representation and algorithm for female red-winged blackbird mate choice. Representationally, one or more features of a territory are represented together as its quality, and the number of females already in a harem is represented separately. The algorithm consists of dividing quality by harem size and selecting the territory with the best ratio. In principle, there are many possible algorithms corresponding to a single solution; successfully identifying a problem and its optimal solution at the computational level does not yield a unique solution at the algorithmic level. However, some algorithms are more parsimonious than others, and a relatively simple decision-making algorithm should generally be preferred over a needlessly complex algorithm that produces the same result. Furthermore, there is no guarantee that the algorithm an animal uses to solve a decision-making problem need actually produce an optimal solution; such algorithms may be satisficing (i.e., guaranteed

to meet some minimal requirements) rather than optimizing (i.e., guaranteed to produce the best possible outcome). These algorithms may, if probed correctly, reveal systematic “bugs”—the equivalent of optical illusions—indicative of the underlying computations involved. In the case of foraging, such algorithms involve decision heuristics like threshold comparison for patch leaving and the 0 – 1 rule for prey selection, as well as more complex rules like *take-the-best* or *elimination by attributes* (Gigerenzer and Goldstein 1996). In some environments (presumably those typical for the species that apply them), these algorithms may perform equivalently to truly optimal algorithms, though there is no guarantee that they will do so in other contexts.

Finally, at the *implementation level* of analysis, we are concerned with the question of which *proximate mechanisms* actually *implement* algorithms like threshold comparison and take-the-best. In principle, there is a many-to-one relationship between implementation and algorithm. The true implementation may involve neuronal or genetic circuits, neuromodulators, hormones, or a complex interplay of these. In fact, a genetic circuit and a neural network may implement the same decision rule—equivalence at the algorithmic level—though their biological details remain entirely distinct. In practice, once we have identified the proximate mechanism by which a decision is made, we can often specify the algorithm, as well.

Decision neuroethology proceeds by considering all three levels of analysis. The computational level identifies the decision-making problem to be solved. The algorithmic level proposes one or more means for computing approximate solutions to the problem, and the implementation level identifies the actual proximate mechanisms by which the algorithm is performed. In practice, the study of algorithms may suggest implementations or vice versa. It is also interesting to consider that implementations themselves may have costs in terms of caloric consumption or ontogenetic complexity, which may be analyzed at the computational level to understand why one implementation is favored over another, or even why some algorithms are favored if they happen to be readily implemented by low-cost mechanisms.

In what follows, we will explore this threefold description as a means of exploring what appear to be striking similarities across taxa in behavioral ecology—similarities, we suggest, that result from convergent evolution at the algorithmic level to common problems posed at the computational level. We examine these ideas through the lenses of two of the most ubiquitous problems in behavioral ecology: patch foraging and social information seeking.

Patch-Leaving Decisions

A well-studied example of an ecological decision is the patch-leaving problem, mathematically analyzed by Charnov (1976) and first tested in a series of experiments in birds by Krebs et al. (1974). This problem considers an animal

foraging in an environment with food items distributed in sparse patches. As the animal forages in a patch, local resources are depleted, and the time required to find a new food item increases, thus reducing the rate of food intake. As a result, animals must balance the benefits of diminishing returns against the costs of searching for new patches. The optimal solution to this problem, the *marginal value theorem* (MVT; see Kacelnik, this volume)—that foragers should abandon patches when the local rate of caloric return falls below the average for the environment as a whole—has been demonstrated to hold in a breathtakingly wide array of species, including worms, insects, fish, rodents, birds, nonhuman primates, and humans (Stephens and Krebs 1986; Stephens et al. 2007).

On the computational and algorithmic levels, this result is unsurprising: most animals must forage, and nutrients are often sparse, so many species face computationally equivalent problems. Still, the replication of a simple decision rule across species with such diverse neuroanatomical organization need not imply conserved proximate mechanisms. Rather, algorithms useful for solving ubiquitous problems like patch leaving and prey selection (see below) are more likely to be products of convergent evolution, primarily because they are robust, require only simple components, and do not require a centralized architecture. Such algorithms are thus more likely to be repeated across taxa.

In a recent experiment investigating the neural basis of patch-leaving decisions, Hayden, Pearson, and Platt (2011) designed a laboratory version of the problem in which monkeys chose between “stay” and “leave” options represented by visual stimuli on a computer monitor. Monkeys received juice rewards of diminishing value for selecting the stay option, simulating the effect of remaining in a depleting food patch, whereas selecting the leave option resulted in a cued delay, simulating the travel time between patches. As in previous studies in other species, the authors found that monkeys readily optimized their patch residence times (Figure 6.2). Monkeys’ patch-leaving decisions also depended systematically on the “travel time” to the next patch, consistent with the MVT. Notably, monkeys slightly, but systematically, remained in patches for longer than predicted by the MVT. As noted elsewhere (Cuthill et al. 1990), however, this finding is precisely the result expected if monkeys’ estimates of reward rate were based on short-run returns, a quantity calculable by simple linear filtering (Anderson and Moore 1979; Bateson and Kacelnik 1996). As a result, this deviation from optimal behavior itself carries information about the underlying decision algorithm, utilizing an average of recent returns that can be linearly updated each trial instead of the full nonlinear computation of reward rate across the entire foraging history.

Simultaneously, Hayden and colleagues recorded the firing patterns of single neurons in the dorsal anterior cingulate cortex (dACC), an area of the macaque and human brain linked to reward monitoring, error signaling, learning, and behavioral control. Neuronal firing rates revealed a strikingly simple implementation of a thresholded decision circuit (Figure 6.3). Single neurons




Figure 6.2 Rhesus macaques forage nearly optimally in a computerized patch-leaving task. Monkeys remain in the patch longer as travel time rises, as predicted by the marginal value theorem (MVT). Each dot indicates a single patch-leaving decision ($n = 2,834$ patch-leaving events) sampled from two individuals. The time at which the monkey chose to leave the patch (Y axis) was defined relative to the beginning of foraging in that patch. Travel time was kept constant in a patch (X axis). Data from both monkeys is shown. Behavior (average is traced by the blue line) closely followed the rate-maximizing leaving time (red line), albeit delayed by 0–2 s.

responded phasically (i.e., with a transient, time-locked change in firing rate) to each decision. Further, phasic neuronal activity increased with each successive decision to stay in a given patch, across multiple actions unfolding over tens of seconds. Finally, firing rates peaked with the decision to abandon a patch and move on to the next. In fact, peak firing not only predicted when monkeys decided to leave the patch, it also differentiated between premature and postponed leave decisions. Furthermore, with increasing “travel time” between patches, the rate of increase in peak firing diminished, in keeping with the longer dwelling times observed behaviorally.

All of these observations are consistent with a firing rate threshold for the leave decision; when neuronal activity in dACC reaches this fixed value, monkeys opt to leave the patch. More specifically, firing rates on patch-leaving trials did not differ statistically across premature or postponed leaving decisions within a given delay condition, further strengthening the conclusion that the threshold is constant for a given neuron. Nonetheless, leaving thresholds for firing rates did increase with delay to the next patch, indicating a flexible control process necessary to reset the threshold to generate optimal behavior for the current environment.

The rise-to-threshold process evident in the decision-related responses in ACC mirrors similar rise-to-threshold processes evident in the activity of neurons in parietal and prefrontal association cortex in monkeys rendering perceptual judgments (Gold and Shadlen 2007). Further, such integrate-to-threshold processes have been theorized to serve as implementations of the sequential probability ratio test (SPRT)—the most efficient solution to this type of binary decision problem (Wald and Wolfowitz 1948; Ratcliff 1978; Ratcliff and McKoon 2008). The observation that a similar process appears to govern both

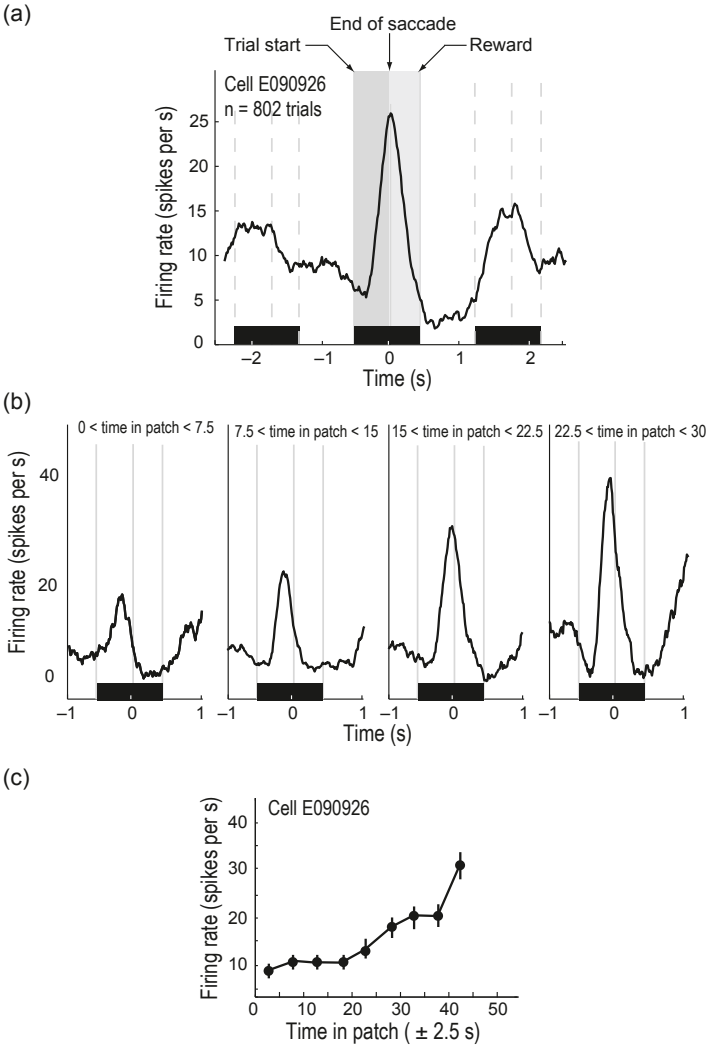


Figure 6.3 The rising value of leaving a patch is represented by single neurons in the macaque frontal cortex. (a) Average reward-aligned peri-stimulus time histograms (PSTHs) are shown for example cell in the anterior cingulate cortex (ACC). Neuronal responses were briefly enhanced around the time of saccades and then fell to a baseline level between trials. Time zero indicates end of saccade, indicating choice. Dark gray box indicates the pre-saccadic epoch; light gray box shows the post-saccadic epoch, and the black rectangle indicates the average duration of the trial. (b) The firing rate during the peri-saccadic period rose with time in patch. Each panel indicates responses selected from one range of patch residence times. (c) Average responses of example neuron occurring in a series of 5-second analysis epochs (gray box in (a)). Firing rates increased as time in patch increased. Error bars represent standard error of the mean.

perceptual decisions and patch-leaving decisions endorses the idea that the brain uses a small suite of common mechanisms to solve diverse problems in multiple domains (Hayden et al. 2011).

Recent evidence also indicates that genes coding for neuromodulatory chemicals may mediate the impact of local environmental conditions on patch-leaving decisions, perhaps by controlling the threshold for patch-leaving decisions. Bendesky and colleagues (2011) compared patch-leaving decisions by a nematode (*Caenorhabditis elegans*) strain from Hawaii with one developed in the laboratory. Hawaiian worms abandoned algae lawns at ten times the rate that the laboratory strains did. The authors also found that differences in patch-leaving threshold among distinct strains resulted from polymorphisms in promoter regions of the tyramine receptor gene (*tyra 3b*), which controls the expression of a G-protein-coupled receptor analogue of vertebrate catecholamine receptors. Thus worms' patch-leaving thresholds are regulated genetically, as opposed to monkeys' use of flexible thresholding, though both obey the MVT. In other words, completely different proximal mechanisms are capable of instantiating the same algorithm to solve a computationally equivalent biological problem.

Intriguingly, the invertebrate catecholamines tyramine and octopamine, which bind to the *tyra3* receptor, are closely related structurally to the vertebrate neuromodulator norepinephrine, which has been hypothesized to regulate exploration/exploitation trade-offs in primates and humans (Aston-Jones and Cohen 2005; Cohen et al. 2007). Here again, a similar design for regulating the patch-leaving threshold is implemented by two separate mechanisms—one genetic, the other cortical—with behaviorally similar results. It is tempting to speculate that individual differences in exploration behavior, an innate tendency to abandon the current behavioral strategy for another potentially more profitable one, might also be mediated in humans through genetic influences on catecholaminergic neuromodulatory systems (Frank et al. 2009).

Social Information Seeking

The ability to select, inhibit, and shift behavior rapidly is particularly important in social species, given the highly dynamic nature of social environments. For example, the act of consuming a food resource is often mutually exclusive with vigilance behavior. This behavioral trade-off is analogous to the explore/exploit decisions made between two foraging sites. Thus the primacy of information in guiding decisions is abundantly evident in the social behavior of primates. Primates have frontally oriented, mobile eyes with a central fovea composed of a high density of cone photoreceptors, and thus are not capable of sampling all regions of the visual field simultaneously. This set of adaptations has led to the evolution of mechanisms which orient the visual system to objects with high information value via overt and covert attention (Moore et al.

2003). In the context of foraging, information is given value as a consequence of the value of the nutrient it may yield, as when an animal evaluates a tree for the presence of ripe fruit. Animals living in complex and dynamic societies can, however, use the same attentional strategies to gather information about others (Klein et al. 2008), including rank (Bovet and Washburn 2003), identity (Parr et al. 2000), group membership (Mahajan et al. 2011), direction of gaze (Lorincz et al. 1999; Ferrari et al. 2000; Deaner and Platt 2003), and emotional state (Sackett 1966). This information must be perceived and evaluated to guide adaptive behaviors, such as abandoning a food resource in the presence of a threatening dominant individual.

Whereas locomotion imposes high energetic costs during foraging, the metabolic costs of information seeking are by contrast quite low. However, information seeking can impose time costs, requiring animals to forego activities such as sleeping, drinking, or eating, which demand postures or behavioral states incompatible with attentive orienting (Figure 6.4). Other costs are social: in the case of rhesus macaques, visual fixation on the face of another individual invites aggression (van Hooft 1967). Furthermore, inappropriate information seeking (e.g., directing attention to low-value information) can result in missed opportunities to gather more useful information elsewhere. These constraints can be considered as part of the computational problem social animals must solve in deciding how to allocate their time and attention in seeking social information.

Thus, as in foraging, the net gains that accrue from information seeking can often outweigh the potential costs, and adaptive decision making depends on the assessment and comparison of these costs and benefits. In an experimental measure of the relative value of different classes of social rewards, Deaner, Khera, and Platt (2005a) demonstrated that thirsty male rhesus macaques will forego a small amount of juice to acquire specific types of social information,



Figure 6.4 Social information is a valuable resource for macaque monkeys. A rhesus macaque on Cayo Santiago assumes a vulnerable posture (left) to drink from a puddle, but periodically interrupts this posture to visually scan the surrounding region for potential threats (right). There are no predators on the island, but aggressive social interactions are commonplace.

such as reproductive signals (i.e., female perinea) or the faces of dominant males, but will not do so for other types of social information, such as the faces of subordinates (Figure 6.5). An alternative measure of value—the duration of time that the monkey chooses to look at the image once it is displayed—shows that monkeys look longest at reproductive signals but quickly avert gaze from both dominant and subordinate faces. Taken together, these two results invite the hypothesis that although both sexual signals and status-related signals contain high information value, there is a high cost associated with an extended period of direct eye contact. In support of this interpretation, a genetic polymorphism in the serotonin system associated with heightened anxiety in humans elicits reduced attention to the faces of other monkeys, greater pupil dilation (a somatic index of elevated autonomic arousal) in response to faces of dominant males, and reduced reward value for viewing the faces of dominant males (Watson et al. 2009) in rhesus macaques.

Patterns of neuronal responses also support the notion that information gathering has value for making decisions. Dopamine neurons, which respond to primary reinforcers (such as nutritive rewards) when unpredicted and to cues that predict them, also encode monkeys' preferences for advance information about impending choices between primary rewards (Bromberg-Martin and Hikosaka 2009). Moreover, the firing of neurons in the lateral intraparietal area (LIP)—a region of visuomotor cortex thought to encode a saliency map of the visual world (Goldberg et al. 2006; Bisley and Goldberg 2010)—varies with the value of social images displayed in the neurons' receptive fields. Neurons in LIP not only encode the value of the juice reward they will gain for orienting to a particular location (Platt and Glimcher 1999), but also the value of the social information they receive for orienting to the same location (Klein et al. 2008). Importantly, social and gustatory value are encoded independently using the same coding scheme, suggesting that LIP plays a role in assigning value to a particular location in space.

In another experiment, a subpopulation of LIP neurons showed increased activity when a centrally positioned monkey face oriented its gaze toward the neuron's receptive field. This result is consistent with the well-known *Posner effect*, in which reaction times toward an eccentric target are reduced when attending a target and increased when attending away from the target (Posner et al. 1980). Both of these studies bear obvious relevance to the phenomenon of *joint attention* (Scaife and Bruner 1975), in which gaze is reflexively drawn in the direction of another's line of sight, an example of socially facilitated information gathering that is severely disrupted in neuropsychiatric disorders such as autism and schizophrenia.

Regions in the visual perceptual cortex in the temporal and occipital lobes are specialized for processing social information in humans and macaques, suggesting adaptation for the rapid assessment of visual social information. Humans and macaques both possess multiple brain regions, identified by functional imaging and confirmed by recordings from single neurons, exquisitely

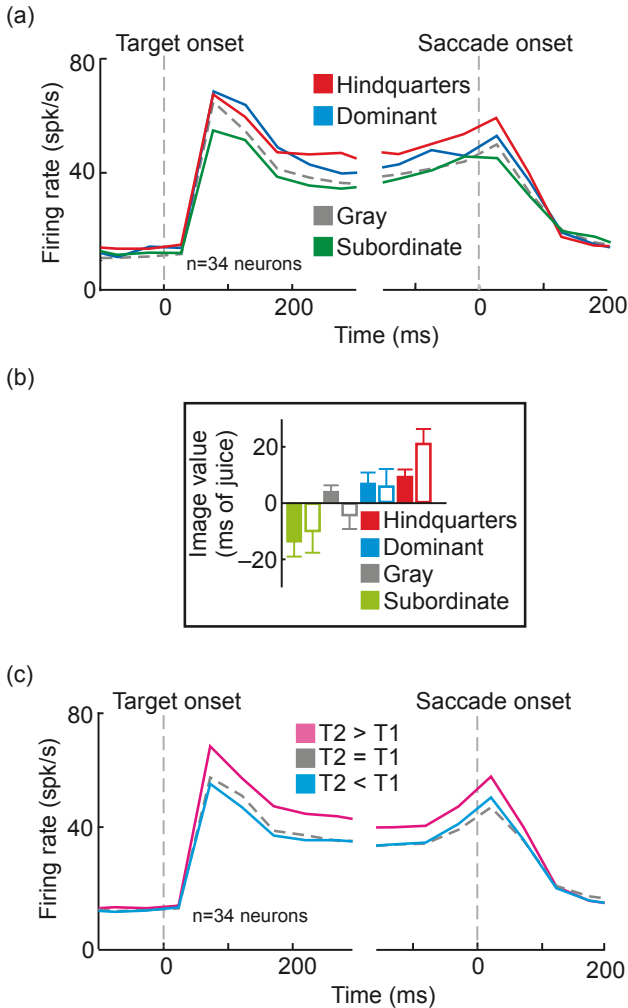


Figure 6.5 The value of social information is signaled by neurons in the macaque visual-orienting system. (a) Average firing rate for 34 LIP neurons plotted against time for all trials in which the subject chose to view the image (T2) in the neuron's response field, separated by image class. (b) Values determined for different image classes for two male monkey subjects (open and closed bars), in ms of fluid delivery time. Positive deflections indicate the subject was willing to forgo fluid to view that image class. Negative deflections indicate the subject required fluid overpayment to choose that image class. The category "hindquarters" refers to the perineal sexual signals of familiar females; "dominant" and "subordinate" to the faces of familiar dominant and subordinate males; and "gray" to a plain gray square matched for size and luminance to the other image classes. (c) Average firing rate of the population for all trials in which the subject chose to view the image (T2) in the neuron's response field, separated by fluid value relative to the nonchosen target (T1).

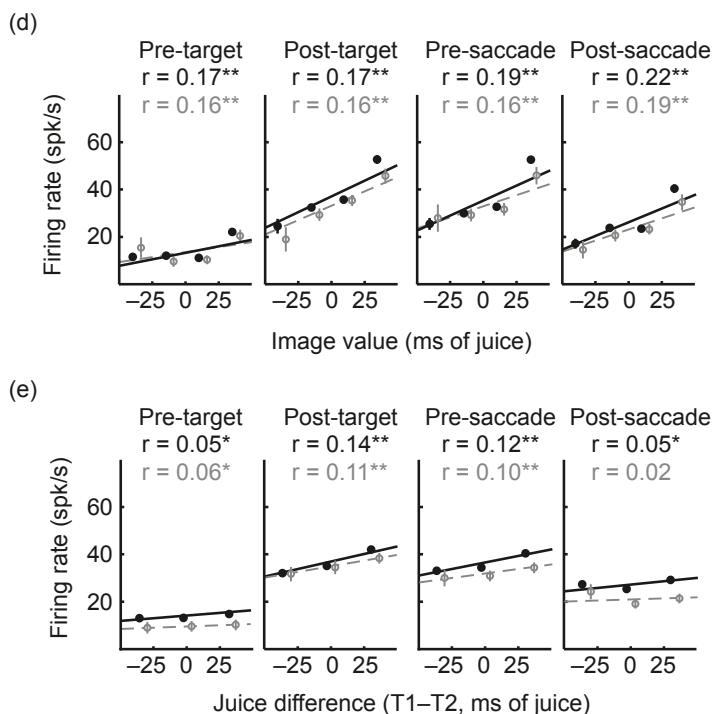


Figure 6.5 (continued) (d) Firing rates plotted as a function of image value in four 200-ms epochs. Black symbols represent regressions performed on all data in which the subject chose to view the image; gray symbols depict the same analysis restricted to trials in which the fluid payoff for choosing T1 was equal to T2. (e) Firing rates plotted as a function of the difference in fluid payoff between T2 and T1. Black symbols represent regression performed on all data in which the subject chose to view the image; gray symbols depict the same analysis restricted to trials in which the image value calculated for that block was greater than -5 and less than 5 ms. Error bars represent standard error of the mean. The data in (c) and (d) were binned for display, but all regressions were performed on raw data. * $p < 0.05$, ** $p < 0.001$.

selective for faces (Kanwisher et al. 1997; Tsao et al. 2003). Neurons in other temporal lobe areas encode head direction (Perrett et al. 1985), face identity (Perrett et al. 1984), or biological motion (Oram and Perrett 1994). These populations of neurons are interconnected and feed forward from the temporal into the frontal lobe, where various modalities and features of external stimuli, including their motivational value, are integrated (Ku et al. 2011). The architecture of this network suggests that frontal decision-making mechanisms may have privileged access to social information in human and nonhuman primates alike.

Despite the presence of dedicated social perception networks in the visual systems of primates, there is little evidence to suggest that the brain regions

involved in decision making, including the ventral tegmentum, basal ganglia, and frontal cortex, contain specialized subnetworks devoted to social decisions. This observation invites the hypothesis that the rise-to-threshold neural mechanism is applicable to social information gathering, as well as to perceptual and foraging decisions. Given the example in which a subordinate monkey places himself in a vulnerable position consuming a food resource, one could imagine that the presence of a potential social threat would induce a rise in firing rate in the ACC, causing the monkey to shift away from his current activity in favor of vigilance or avoidance behavior. More generally, we conjecture that a common suite of algorithms subserves both “social” and “nonsocial” decision processes, with the former differentiated from the latter by the large number of specialized structures used for detecting and processing information related to conspecifics.

Discussion

The study of ecological decision making suggests the possibility that natural selection has favored a set of simple, repeated design patterns: basic circuit elements capable of being implemented by many biological configurations. Instead of forming a single unified system for decisions, these local circuits might be capable of functioning independently for specialized subclasses of action planning and selection, as well as being recruited across regions for more complicated behaviors. Just as Gigerenzer has suggested that human decision processes draw on a well-stocked “adaptive toolbox” filled with inexpensive, approximate heuristics, so evolution appears to favor repeating algorithms, despite disparate implementations (Gigerenzer and Selten 2001).

In this view, the study of a decision-making problem like patch leaving is valuable not only because it is ubiquitous, but also because the algorithm used to solve it—comparison of local returns to a fixed threshold—represents one of the simplest forms of a single-input/single-output control system, in which the controller implements a binary threshold operation (Brogan 1985). Any neural system capable of implementing such a circuit is likewise applicable to an extremely wide class of problems, one that extends far beyond foraging. In fact, a more general version of such a system, the Kalman filter, is known to be an optimal solution to the problem of predicting returns under fairly general assumptions, and thus for fine-tuning behavior in response to changing environmental conditions (Anderson and Moore 1979; Brogan 1985). That such a system is linear, that it requires only a simple architecture, and that it is robust against noise all lead us to expect that it will not only appear in a wide diversity of species, but that it may be repeated and repurposed within a single brain to solve seemingly unrelated problems. In fact, such observations motivate a neural engineering viewpoint in which the unique classes of problems faced by an

organism become primary, followed by the algorithms used to solve them, and only last their specific neural implementations (Marr 1982).

This does not mean, however, that comparative biology or neurophysiology become irrelevant. On the contrary, this evolutionary viewpoint suggests that algorithms implemented in more conservative nervous systems are more applicable than we might have thought. Just as in vision or olfaction, the insights gained from studying flies or worms suggest possibilities at the algorithmic level in birds or primates. Though the details may differ as organisms become capable of more generalized and flexible behaviors, the same simple biological components, coupled-like circuit elements, may likewise be expected to give rise to startlingly sophisticated generalizations (Brogan 1985).

Many models of decision making, particularly those derived from economics, describe the decision process as a linear sequence of first estimating abstract utility of several possible outcomes or behavioral plans, directly comparing these utilities, and finally selecting the goal or behavior associated with the highest utility (Glimcher 2004; Glimcher et al. 2005; Sugrue et al. 2005; Lee 2006; Padoa-Schioppa 2007). In some cases, such models leave the exact nature of the abstract utility undefined. Here, however, we have presented a “bottom-up” as opposed to “top-down” perspective, in which simple, reusable decision rules substitute across taxa for what are often considered outputs of a single decision-making system. We argue that these design patterns, implemented by diverse suites of neural hardware, should nonetheless prove ubiquitous on evolutionary grounds, and that their simplicity and robustness should favor them both for convergent evolution and conservation within taxa. Such claims represent a new opportunity for both systems theory and comparative biology, since the view of decision systems as evolving primarily to solve ecological problems demands renewed interest in both engineering disciplines and animal behavior. Indeed, the search for reusable design patterns in neural systems may provide a unifying framework for biological decision making in much the same way it has for vision and motor control.